AFRL-VA-WP-TP-2002-307

RELATING AUTONOMY TO A TASK - CAN IT BE DONE?

Bruce T. Clough

APRIL 2002

Approved for public release; distribution is unlimited.

20020912 085

AIR VEHICLES DIRECTORATE
AIR FORCE RESEARCH LABORATORY
AIR FORCE MATERIEL COMMAND
WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7542

REPORT DOCUMENTATION PAGE

1. REPORT DATE *(DD-MM-YY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
April 2002	Conference Paper Preprint	03/01/01 – 03/01/02

4. TITLE AND SUBTITLE		
RELATING AUTONOMY TO A TASK - CAN IT BE DONE?	**5a. CONTRACT NUMBER** IN-HOUSE	
	5b. GRANT NUMBER	
	5c. PROGRAM ELEMENT NUMBER 62201F	

6. AUTHOR(S)	
Bruce T. Clough	**5d. PROJECT NUMBER** 232F
	5e. TASK NUMBER FE
	5f. WORK UNIT NUMBER PA

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
Control Systems Development and Application Branch (VACC) Control Sciences Division Air Vehicles Directorate Air Force Research Laboratory, Air Force Materiel Command Wright-Patterson Air Force Base, OH 45433-7542	AFRL-VA-WP-TP-2002-307

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSORING/MONITORING AGENCY ACRONYM(S)
AIR VEHICLES DIRECTORATE AIR FORCE RESEARCH LABORATORY AIR FORCE MATERIEL COMMAND WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7542	AFRL/VACC
	11. SPONSORING/MONITORING AGENCY REPORT NUMBER(S) AFRL-VA-WP-TP-2002-307

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited.

14. ABSTRACT *(Maximum 200 Words)*
There are currently countless methods touting being the best at autonomous control. Whether it is emergent behavior or dynamic programming, fuzzy logic or Bayesian Belief Networks, each technology has its phalanx of zealots that proclaim it to be superior for the job and the rest rotten, or at least trivialized. In this paper, we step back and look at the big picture in autonomous control, developing the idea that what will drive the choice of the particular technology is the task, or tasks, that the autonomous system must do, and not the merits of any particular technology by itself. The optimal technology for any task in inexorably intertwined with the task – they cannot be separated.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT:	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON (Monitor)
a. REPORT Unclassified	**b. ABSTRACT** Unclassified	**c. THIS PAGE** Unclassified	SAR	16	Bruce T. Clough **19b. TELEPHONE NUMBER** *(Include Area Code)* (937) 255-2831

RELATING AUTONOMY TO A TASK – CAN IT BE DONE?

Bruce T. Clough, Technical Area Leader: Control Automation
Air Force Research Laboratory
Wright-Patterson AFB, OH

Abstract

There are currently countless methods touting being "the best" at autonomous control. Whether it is emergent behavior or dynamic programming, fuzzy logic or Bayesian Belief Networks, each technology has its phalanx of zealots that proclaim it to be superior for the job and the rest rotten, or at least trivialized. In this paper we step back and look at the "big picture" in autonomous control, developing the idea that what will drive the choice of the particular technology is the task, or tasks, that the autonomous system must do, and not the merits of any particular technology by itself. The "Optimal" technology for any task in inexorably intertwined with the task – they cannot be separated. In the world of robotics, one would say the technology must be imbedded with the task to make any sense. We develop the idea that tasks and techniques are linked in several manners: deliberate versus reactive, distributed versus centralized, functional versus semantic, and local versus global information requirements. We espouse the idea that no technology is "bad", just mis-matched with the task required. We also discuss the idea that for any task, the amount of information to do it correctly is invariant, distributed amongst communications, sensing, and organic databases, the exact ratio up to the designer. As implementers of autonomous control technology, we are not enamored with any particular technique or architecture, we just want the technology that's right for the task, technology that can be verified and validated with minimal time, effort, and cost. That's the bottom line in bringing autonomous control to the masses.

Introduction

About ten years ago I attended a controls conference about the practical application of modern control techniques. In one of the sessions an argument broke out between the supporters of several different techniques, each stating that it was the best to use. It got very heated, to the point that I thought fists were going to fly. Personally, I thought this was extremely

---

entertaining, since as a veteran engineer of several advanced engineering controls programs I knew that none of their techniques worked exceptionally well in the real world. They all had benefits and drawbacks, more drawbacks, it seemed, than simple PID (proportional, integral, derivative) controllers!

I now see this in autonomous vehicle control. I specialize on unmanned aerial vehicles (UAVs), transitioning autonomous control systems from the laboratory into flight vehicles. Over the last ten years our organization has been developing control laws for UAVs, first the inner (vehicle control) loops, then the outer (flight management) loops. In doing so we have bumped into a wide variety of autonomous system design techniques. Each vendor assures us that their techniques are the best for our application, and the other's techniques are not as good. Again, observation says that nobody's technique is the best for all cases, all tasks. When we build our autonomous systems we tend to use a mix of techniques, matching the task to the technique. We've never really formally stated why this is so, it just "feels good" to us. This paper explores why certain techniques work well for some tasks, and others don't. What we have found out is that one needs to match the autonomous control method to the task, with the control technique's strengths feeding into the tasks specific characteristics. The autonomous control system designer must be ready to apply a broad range of techniques to solve his/her challenges.

Note that this paper contains no equations, no axioms, and no proofs. The actual world of autonomous control is highly non-linear, highly discontinuous, chaotic, and non-deterministic. We leave it to the reader to try and develop formal proofs!

Background

Figure 1 shows our goals in autonomous control over the next decade or so, split into three phases. The crux of the matter is we are planning on developing the technology to do any USAF mission autonomously. Whether or not that will actually occur will be up to the senior DOD leadership and the politicians, not because the technology doesn't exist.

In order to do this we are investigating what has to occur to take those tasks, which up to this point were done by humans, and allow the machine to accomplish

Autonomous Control
Future Capability

Near Term   Mid Term   Far Term

Control Attributes

Control Attributes
• Collaborate With Man
• As Reliable As Man
• Flexible Direction
• Intelligent Packs
• Relaxed Ops
"Monitored"

Control Attributes
• Co-Habitate With Man
• More Reliable Than Man
• Autonomous Direction
• Intelligent Swarms
• Seamless Ops
"Autonomous"

"As autonomous as needed, as interactive as desired"

Figure 1: Autonomous Control Goals

them. We don't know if in the specific application that the human will allow it to do that task, but since we don't know that he/she won't, we have to plan accordingly and deliver the capability of doing the task to our users and developers. We will not be a limiting factor; we will allow the user to be "as autonomous as needed, as interactive as desired" [1]. We have put technology plans in place to make this so, and in the process have relied on a number of different techniques. This paper feeds on the lessons learned by us as we've stive to meet the goals laid out in Figure 1.

Our Use Of "Autonomous" And The Difference Between Autonomous and Automatic (our definition)

But before we start our discussion on techniques, let us quickly review the word autonomous. We need to do this to insure the reader is working in the same reference system as we are. "Autonomous" in this paper means "the system has the freedom to make the choice". It does not mean that the UAV isn't in communications with humans as some autonomy definitions state - it could be continuously chatting with humans. It's just when the time comes to make a decision, it's made by the UAV. It has the human's proxy. Autonomy is our ability to give the UAV our proxy.

Also, many people don't realize that there is a significant difference between the words autonomous and automatic. Many news and trade articles use these words interchangeably. Automatic means that a system will do exactly as programmed, it has no choice. Autonomous means that a system has a choice to make free of outside influence, i.e., an autonomous system has free will. For instance let's compare functions of an autopilot and an autonomous guidance system:

• Autopilot: Stay on course chosen.

• Autonomous Guidance: Decide which course to take, and then stay on it.

For instance, a cruise missile is not autonomous, but automatic since all choices have been made prior to launch.

Relating Autonomy To The Task – First Step: Determining Task Attributes

Remember above, where we said the goal was to relate the technique to the method, feeding into the techniques strengths? In order to do this we need to determine the task attributes. Figure 2 is our "cut" at what these characteristics are which seem to influence our choice of technique. Let's spend a few words on each:

2

Task Characteristics
Is It...   Or...

Is It...	Or...
Well Defined	Ill Defined
Time Invariant	Time Dependent
Simple	Complex
Reactive	Planned
Centralized	Distributed
Local Info	Global Info
A priori Info	No A Priori Info
Stand Alone	Integrated
Critical	Supporting
Lethal	Non-Lethal

Figure 2: Task Characteristics

Well Defined Versus Ill Defined

Have we thought through all aspects of the task, or is it a vague notion? How well have we set it forth, or how well can we? It's the difference between "lay out the red and blue sticks, placing the blue sticks into the center holes and putting the red sticks in the circumferential holes" and "oh, just go put the Tinkertoys$_{tm}$ together".

Time Invariant Versus Time Dependent

Does it matter when it is done, or does it have to be done at a specific instant? Is it time critical?

Simple Versus Complex

Is it easy to comprehend and think through, or is it very difficult to articulate and dissect all the possible ramifications?

Reactive Versus Planned

Is the task a reflex to external events, or is it following a plan to make others react to it?

Centralized Versus Distributed

Can it be confined to one location, or must it occur at multiple locations simultaneously?

Local Versus Global Information

Does it take a broad general knowledge of what is going on - does it require the "big picture", or does it just require information of local circumstances?

A Priori Versus Non-A Priori Information

Is a lot of information required beforehand, or none? For instance, the a priori information requirements between "take everything out in the Kill Box" are a bit different than "Take everything out in the Kill Box except non-combatants, friendlies, and left-handed people".

Stand Alone Versus Integrated

Is the task done by itself, or is it just a subtask in the greater scheme of things that has to be properly sequenced and coordinated?
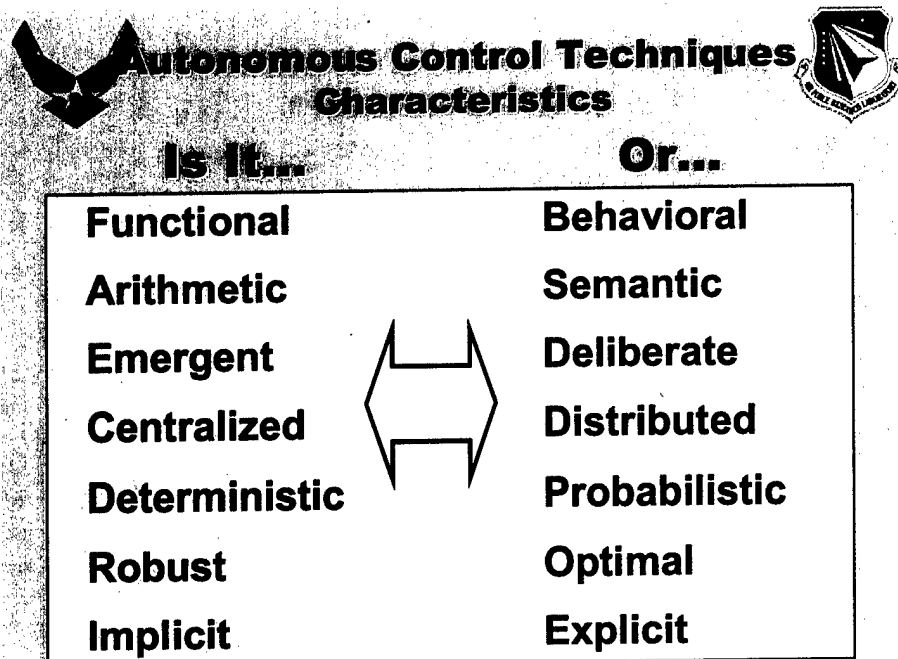
American Institute of Aeronautics and Astronautics

Figure 3: Autonomous Control Technique Characteristics

Relative Importance To Mission: Critical Versus Supporting

Is the task the fulcrum to which the mission rotates around, or is it ancillary to the primary goal? Can the mission go forward even if the task never completes?

Lethal Versus Non-Lethal

This last one seems a bit different than the rest, but it is extremely important to put adequate controls on any system, carbon or silicon based, that could employ lethal force in it's task.

There are other possible attributes, such as how integrated the task is with humans, how much resources it uses, and is it an individual or team pursuit; however, we feel that those are derived attributes of the above.

Step 2: Determining Autonomous Control Technique Attributes

We can go ahead and do the same type of categorization for the autonomous control techniques we use. Figure 3 shows the ones we've identified. As with the task characteristics, we will spend a few words on each:

Functional Versus Behavioral

Does the technique treat the control process as a function that must be evaluated to arrive at an answer, a way to get a closed-form solution, or behaviors that are executed with the solution a byproduct of their operation?

Arithmetic Versus Semantic

Does the technique rely on mathematical expressions and numbers to encapsulate the control and knowledge, or does it rely on words and symbols? [*Note 1*]

Emergent Versus Deliberate

Does the technique rely on reactive functions or behaviors, the interactions of which emerges the control desired, or does it develop and execute overt plans?

Centralized Versus Distributed

Does the technique assume control/decisions occur at a central location, or is it dispersed amongst multiple agents?

Deterministic Versus Probabilistic

Does the technique assume determinism for the inputs and outputs (deterministic mapping), or does it result in

4

American Institute of Aeronautics and Astronautics

probabilistic distribution of outputs even with deterministic inputs? [*Note 2*]

Robust Versus Optimal

Does the technique always strive to give the best performance relative to some metric, or does it give acceptable performance over wider environments? [Note 3]

Implicit Versus Explicit

Does it rely on models/plans explicitly programmed into the code, or does it carry those implicitly programmed through code development? For instance, in the control theoretic world a PID controller carries the model implicitly since the models were used in development and tuning, but they are not directly coded into the controller vis-a-vis a controller based on a kalman filter with internal models.

Attribute Triage

Fine. We've spent the last several pages developing sets of metrics describing tasks and autonomous control techniques. Our goal is to compare the two in order to come up with some insight in linking technique to task, so how do we compare the ten task and seven technique metrics?

Well, first of all we make the observation that what really matters are the task characteristics since that is what we must accomplish. The techniques are just used to accomplish the tasks.

Second of all, ten attributes are too many, especially for a technical paper that is developing a small-dimensional metric for presentation at a symposium, or it's other intended application, explaining to management why entertaining ways of doing things make sense economically and technically [*Note 4*]

Task Metric Reductions

The first grouping done was to put critical/supporting, stand alone/integrated, time invariant/time dependent, and lethal/non-lethal into the planned/reactive "bucket". If it's critical, it must get done, therefore one wants to keep track of it – make a deliberate action to do it. If the task is integrated then it must be done at a certain time in a certain order, there again deliberate actions are required. If it is time dependent, this again implies order, implying a plan, implying deliberate actions. If it is lethal, we want to know that it executes on the enemy, rather than on friendlies or non-combatants – we want the release of lethal force to be very deliberate.

Next we integrated the infocentric characteristics together. The a priori information characteristics were integrated with local versus global info required. Reasoning: if no a priori information is needed on a task, then it probably is dependent on local, not global, information [*Note 5*].

Finally, we let the centralized/distributed and well/ill defined characteristics stand as is.

Grouping of Autonomous Control Techniques.

Reducing these characteristics, grouping them in categories similar to the tasks, was done along the same lines as the tasks.

First of all, we grouped emergent/deliberate, deterministic/probabilistic, and robust/optimal into the same category, renaming it reactive/deliberate. Emergent systems are reactive anyway, and if we want to enforce determinism we have to be deliberate. Optimality means we have to know what we're optimal to – we have to have a model, which means we have a plan, which means we're deliberate.

Secondly, we grouped functional/behavioral and

Technique	Planning Requirements		Info Requirements		Location		Construction	
	Deliberate	Reactive	Local	Global	Central	Distributed	Functional	Semantic
Dynamic Programming	D			G	C		F	
Emergent Behavior		R	L			D		S
Bayesian Belief Networks	D		L			D	F	
Traditional AI Rules Base	D			G	C			S

Table 1: Quick Comparison Of Technologies Versus Task Characteristics

Task	Planning Requirements		Info Requirements		Location		Construction	
	Deliberate	Reactive	Local	Global	Central	Distributed	Functional	Semantic
Weapons Release	D		L		C			S
Formation Flight		R	L			D		S
Aerial Refueling	D		L		C		F	
Landing	D		L		C		F	

Table 2: Comparison Of Tasks Against Task Characteristics

arithmetic/semantic into one group called functional/behavioral. This was pretty easy to do since functions tend to be arithmetic in nature, and behaviors tend to be best represented via semantic expressions.

We left central/distributed and implicit/explicit in their own groupings since it made sense to do that, especially considering the next few paragraphs.

Getting The Tasks And Autonomy Metrics To Line Up

Doing the above almost got us to our end goal of common metrics required for linking tasks to autonomous control techniques. We managed to develop a reactive/deliberate characteristic in both places, along with a central/distributed one.

For the other two, we noted that well-defined systems usually can be defined functionally. The converse is also true – it's hard to functionally describe some task that is ill defined whereas semantics work well on ill-defined tasks (which is a reason that fuzzy logic works well in categorizing ill defined things [3]). We decided to keep the definitions as functional and semantic. Finally, we relate the idea of local and global info to implicit and explicit internalizations by noting that tasks that rely on global info do that because they have explicit internalization which requires it where as implicit systems tend to rely on local sensing. This leaves us with the four characteristics:

1. Planning Requirements (Deliberate/Reactive)
2. Info Requirements (Local/Global) [Note 6]
3. Control Locality (Central/Distributed)
4. Algorithm Construction (Functional/Semantic)

Ready To Compare – Bring On The Techniques!

Shall we try to categorize based on the metrics? Table 1 is a quick comparison using several popular techniques one uses to build autonomous control systems. Now before some zealot runs off and defends his/her favorite technique let me point out that this is a simple "yes/no" comparison. Sure, there are exceptions to every rule, but in general the techniques in Table 1 have those characteristics, or that's the way engineers have been using the techniques [Note 7].

But this is only half the story. One has to categorize the task also. Table 2 shows the categorization of several different UAV tasks. We did not go in depth on how these were arrived at due to length constraints; however, the reader is encouraged to contact the author to discuss the characterizations in this document.

Comparing Tables 1 and 2 we note that RLDS appears once in each, implying that emergent behavior might be looked at for formation flight. This isn't news to bio-inspired controls developers who have been pushing this way of looking at the task for a number of years. It also isn't any news to human pilots who have been using it for years to fly formations.

One also notes in tasks' planning requirements column that there are more D's than R's. That's because we are working with system designed for war, and despite it's emergent characteristics, war is a very deliberate human act!

Task Versus Technique – Rules Of Thumb

From staring at Tables 1 & 2 we can develop some rules of thumb, somewhat similar to Myers-Briggs$_{TM}$ deployments for humans, on what different

American Institute of Aeronautics and Astronautics

characteristic groups mean to use in the real world [4]. (Note: the "X" is the "don't care" symbol in the below)

- DGXX, XGDX – bandwidth hogs: anything that plans a lot, or is distributed, *and* relies on global data needs lots of information.
- RLXX – Not the first pick for safety critical tasks. Although safety might be argued reactive impulses in humans (pull hand from burners) one can argue safety needs to be overt (don't put your hands on the burners, dummy!)
- DXXX – Overkill for RXXX tasks. Planning for reactive tasks is unneeded, and inefficient.
- XXCF – Normal "Control Theoretic". This is where most of the techniques near and dear to control theoretic folks show up. This also is where a lot of the flight management tasks show up, so maybe this is telling us something?
- DXCX – Normal safety critical task description. We need to be overt to ensure safety, and we need to make sure the human operator has central control if authorizing weapons release.
- DLXX – Will stress onboard databases, whereas RGXX eliminates most of the need for databases. Planning functions based on local data implies a knowledge base onboard to take care of the global implications. Reactive systems receiving global data need no database since they are acting on current data (no plans to track).

These are just a few rules that we have noted. In general, one is urged to plot both the technique and task making sure that they make sense. Also note that the above is not absolute. There may be task/technique pairings that don't line up with our characteristics, but may be perfectly valid to use. As with all real things in the real world, our observations are not perfect, nor do we expect them to be [*Note 8*]!

Task Complexity Not in Here...

Hey wait - aren't we missing something here? Where did the complexity characteristic go in the task description? They went nowhere. This is experience creeping in. We have noticed that the above factors drive the technique used, not the complexity. For instance, space flight, a complex task, works almost entirely off of very functional, deliberate algorithms. While building an ant colony, another complex task, is done entirely by emergent behaviors. All this is conspiring to tell us that there is something else going on here besides just raw complexity – other factors are at work.

Information And Task – Directly Linked

Before we summarize, we'd like to discuss one more ancillary point about information and the task.

The amount of information required to accomplish a task is constant.

Basically, one needs to know so much to do something. This is kind of a common sense point, but one that is lost with most system designers [*Note 9*]. The information required to do a task in inexorably linked to the task. How you get that information is not, but you only have three choices:

- Sensing It
- Someone Telling You (communication)
- Knowing It Already (internal storage)

The cost trade off isn't how much information is required, but the best way to get it at the point of use. Control designers need to ensure that they identify information requirements for cost trades [6].

Summary

By now the reader should have realized that there are valid reasons to link the autonomy development methods to the tasks that need developed. The reader should also take away that there is no perfect method for autonomous control. In fact, from what we know now it's hard to say that there are any particularly good method general-purpose techniques! All have their places to be used, and to be avoided – and those hinge on the task! There are no "Holy Grails" for autonomy development, just a variety of techniques that the designer needs to have in his/her toolkit.

For folks with my background this may be hard to understand, but they have to realize that autonomous control is not just about applying control theory, but is actually the intersection (in a Venn Diagram sort of way) between control theory, psychology, cognitive science, computer science, operations research, biology, and artificial life (apologize to any research disciplines I left out) as shown in Figure 4.

It is important to line up the task with the technique, saving time, money and frustration. This paper is a first look at linking tasks to the autonomous technique required. We have linked task and technique in four areas:

- Planning Requirements
- Information Requirements
- Control Location
- Algorithm Construction

7

Figure 4: What Is The Discipline Of Autonomy?

From this we can generalize attributes, and point out, in general, when mismatches between task and technique occur. It's not perfect, but then again, we live, and develop systems, in the real world where nothing else is perfect either!

In the future we will continue to use our knowledge of autonomy building techniques, not to determine which one is best, but to confront the real challenges faced by us UAV autonomous control types [5]:

- Replicating Pilot Decision Capability
- Increasing UAV Safety
- Ensuring UAV Affordability

These are the actual challenges, arguing about techniques is wasted bandwidth! This is one place where the results are more important than the process.

Acknowledgments

References

1. Quote from James McDowell, AFRL/VACC, on how the user wants us to develop autonomous control for UAVs. 1999.

2. Clough, B. *UAV Swarming? So What Are Those Swarms, What Are The Implications, And How Do We Handle Them?* To be presented at the AUVSI Unmanned Systems 2002. Orlando, FL July 2002.

3. Clough, B., et al. *Aircraft Luggage Vulnerability Model Final Report.* WL-TR-96-3090. May 1996

4. Myers-Briggs Type Indicators (MBTI). Consulting Psychologists Press Inc. www.cpp-db.com 2000.

5. Clough, B. Unmanned Aerial Vehicles: *Autonomous Control Challenges, A Researcher's Perspective.* World Aviation Congress 2000, San Diego, CA October 2000.

6. From personal discussions with Alan Schultz of the Naval Center for Artificial Intelligence. He assured me that the idea wasn't new. He was unsure of the exact reference at the time, and we've not been able to find it to this date.

Notes

1. I suppose I could have discussed it as monotonic/non-monotonic, continuous or discrete, but those are already assuming an arithmetic solution just from the wording...

2. I was going to use the word "non-deterministic" but I chose "probabilistic" instead. The "Non-D" word tends to scare us control theoretic background folks until we realize that the real world is always Non-D!

3. Oh I suppose one could argue that you can have robust optimal techniques, especially

8

since we've spent so much effort over the last twenty years developing them (I'm from a control theoretic background so I can see the questions coming). The fact is if the system is optimized, then it can't be robust almost by definition. This is especially true of linear systems. For instance, there is a direct relationship between the quality factor "Q" of a filter and the bandwidth it will work over. High Q filters are very optimal at eliminating a frequency, but they are intolerant of frequency deviations on the signal to filter. Low Q filters don't filter as much, but they work over wider bandwidths. Another way of looking at this (inspired by an event while writing) is that an extension cord optimized to reduce the length of wire from outlet to laptop PC is not robust to a 3-year old daughter tripping over it, but one with 6 ft of extra cord is! Survivability requires us to back off on performance.

4. In fact, when presenting ideas we need to keep the dimensionality down to three since that's all humans can visualize. We are not going to make it in this paper, though...

5. Or it had better be. Sending an agent into a situation where no a priori information is given, but global information is required is putting a big burden on one's communication systems! Bandwidth is finite!

6. I guess one could argue that local/global information requirements are wrapped up in the planning requirements and/or the control location. We break it out separately since we believe: a) That autonomous control systems are above all else, infocentric, and what they do hinge on the correct receipt, processing, and dispensation of information, and b) Bandwidth is not free, nor is it unlimited. It is a precious resource, so we need to develop and nurture techniques that conserve it.

7. As the reader might have guessed, some have joked that what we've come up with is a Myers-Briggs$_{TM}$ test for control theories! Actually, not that far off, and possibly a topic for another paper, when one considers we are designing systems to replace systems measured using Myers-Briggs$_{TM}$........

8. For instance, virtual leader techniques, which rely on the fact that everyone knows what the plan is, and knows what the others are supposed to do, therefore no communications are needed, would come in at DLDF, which would indicate that they are somewhat robust to global information requirements. However, these systems live in a non-static world, which usually dooms anything that assumes static worlds (such as virtual leader systems). Can you tell the author is not enamored with virtual leader systems?...

9. This truly is a paper in it's own right, and we've noted that we need to write it.